

Anchor points for genome alignment based on *Filtered Spaced Word Matches*

Chris-André Leimeister¹, Thomas Dencker¹, and Burkhard
Morgenstern^{1,2}

¹ University of Göttingen, Department of Bioinformatics,
Goldschmidtstr. 1, 37077 Göttingen, Germany

²University of Göttingen, Center for Computational Sciences,
Goldschmidtstr. 7, 37077 Göttingen, Germany

March 28, 2017

Abstract

Alignment of large genomic sequences is a fundamental task in computational genome analysis. Most methods for genomic alignment use high-scoring local alignments as *anchor points* to reduce the search space of the alignment procedure. Speed and quality of these methods therefore depend on the underlying anchor points. Herein, we propose to use *Filtered Spaced Word Matches* to calculate anchor points for genome alignment. To evaluate this approach, we used these anchor points in the widely used alignment pipeline *Mugsy*. For distantly related sequence sets, we could substantially improve the quality of alignments produced by *Mugsy*.

Introduction

Sequence comparison is one of the most fundamental tasks in computational biology. Here, a basic task is to *align* two or several DNA or protein sequences – either *globally*, over their entire length, or *locally*, by restricting the alignment to a single region of homology. Standard approaches to sequence alignment assume that the input sequences derived from a common ancestral sequence, and that evolutionary events are limited to substitutions, insertions and deletions of single residues or small sequence segments. In this case, sequence homologies can be represented by *global sequence alignments*,

that is by inserting *gap characters* into the sequences such that evolutionarily related sequence positions are arranged on top of each other. Under most scoring schemes, calculating an *optimal* alignment of two sequences takes time proportional to the product of their lengths and is therefore limited to rather short sequences [42, 47, 24, 37, 22].

With the rapidly increasing number of partially or fully sequenced genomes, alignment of *genomic* sequences has become an important field of research in bioinformatics, see [23] for a recent review and evaluation of some of the most popular approaches. Here, the first challenge is the sheer size of the input sequence that makes it impossible to use traditional algorithms with quadratic run time. The second challenge is that related genomes often share *multiple* regions of local sequence homology, interrupted by non-conserved parts of the sequence where no significant similarities can be detected. This means that neither *global* nor *local* alignment methods can properly represent the homologies between whole genomes. Finally, evolutionary events such as duplications and large-scale rearrangements must be taken into account. Since it is not possible, in general, to represent homologies among genomes in one single alignment, advanced genome aligners return alignments of so-called *Locally Collinear Blocks*, *i.e.* blocks of segments of the input sequences that contain the same genes in the same relative order.

Since the late Nineteen Nineties, major efforts have been made to address the problem of genome alignment, and many approaches have been published. One of the first multiple-alignment programs that was applied to genomic sequences was *DIALIGN* [38, 40]. This program composes multiple alignments from chains of local pairwise alignments, and it does not penalize gaps; it is therefore able to align sequences where local homologies are separated by long non-homologous segments. The program has been applied, for example, to identify small non-coding functional elements in genomic sequences [3, 12]. However, the program was initially not designed for large genomic sequences, and it is limited to sequences up to around 10 *kb*. Moreover, *DIALIGN* is not able to deal with duplications, rearrangements and homologies on inverse strands of genomes.

To align longer sequences, most programs for genomic alignment rely on some sort of *anchoring* [30, 39]. In a first step, they use a fast method for local alignment to identify high-scoring local homologies, so-called *anchor points*. Next, *chains* of such local alignments are calculated and, finally, sequence segments between the chained high-scoring local alignments are aligned with a slower but more sensitive alignment method. For multiple sequence sets, anchor points can be defined either between pairs of sequences or between several or all of the input sequences. A pioneering tool to find

anchor points for genomic alignment was *MUMmer* [18]; the current version of the program [32] is considered the state-of-the-art in alignment anchoring. *MUMmer* uses *maximal unique matches* as pairwise anchor points to align genomic sequences or protein sequences. By contrast, *MGA* [28] is a tool for *multiple* alignment of genomic sequences that uses *maximal exact matches* between *all* sequences within a given sequence set. Both *MUMmer* and *MGA* use *suffix trees* [31] to rapidly identify pairs or blocks of identical words, one word from each of the sequences, that are then used as anchor points. Both programs are able to align entire bacterial genomes, *MUMmer* was also used in the *A. thaliana* genome project [48]. However, since the probability of homologous exact matches rapidly decreases with increasing divergence, they are most useful to compare closely related genomes, such as different strains of *E. coli*.

Other approaches to genome alignment are *OWEN* [43], *AVID* [7], *MAVID* [8], *LAGAN and Multi-LAGAN* [10], *CHAOS/DIALIGN* [9], the *VISTA genome pipeline* [21], *TBA* [5] and *Mauve* [15], see [19, 4] for a review. All of these methods are based on alignment anchoring, and most of them are able to deal with duplications and genome rearrangements. Some methods for genomic alignments are based on *statistical* properties of the sequences [6, 15]. Other methods are based on *graphs*, for example on *A-Bruijn graphs* [45] or on *cactus graphs* [44]. A further development of *Mauve* called *progressiveMauve* uses palindromic *spaced seeds* instead of exact word matches as anchor points [16]. That is, for a given binary pattern of length ℓ representing *match* and *don't-care* positions, one searches for a set of ℓ -mers, one ℓ -mer from each of the input sequences, such that all ℓ -mers have matching nucleotides at the *match* positions. At the *don't-care* positions, mismatches are allowed. Palindromic patterns are used to cover both strands of the input sequences. Spaced seeds are used in database searching [36, 17] and alignment-free sequence comparison [33] since they have been shown to lead to better results than contiguous word matches.

Mugsy [2] is a popular software pipeline for multiple whole-genome alignment. In a first step, this program uses *Nucmer* [32] to construct all pairwise alignments of the input sequences. *Nucmer*, in turn, uses *MUMmer* to find exact unique word matches which are used as alignment anchor points. An *alignment graph* is constructed from these pairwise alignments using the *SeqAn* software [20], and *Locally Collinear Blocks* are constructed. Finally, a multiple alignment is calculated using *SeqAn::TCoffee* [46]. *Mugsy* has been designed to align closely related genomes, such as different strains of a bacterium. Here, it produces alignments of high quality. On more distantly related genomes, however, the program is often outperformed by other mul-

tiple genome aligners [23].

Finding *anchor points* is the most important step in whole-genome sequence alignment. Here, a trade-off between *speed*, *sensitivity* and *precision* is necessary. A sufficient number of anchor points is required in order to reduce the search space and thereby the run time for the subsequent, more sensitive alignment routine. Wrongly chosen anchor points, on the other hand, can substantially deteriorate the quality of the final output alignment. If spurious similarities are used as anchor points, this not only results in non-homologous parts of the sequences being aligned. Wrong anchor points may also prevent the program from aligning biologically relevant, true homologies since aligning them may be incompatible with the selected anchors. Also, if the number of anchor points is too large, finding optimal chains of anchor points can become computationally expensive.

In this paper, we propose a novel algorithm to find pairwise anchor points for genomic alignments that is based on the *Filtered Spaced Word Matches (FSWM)* idea that we previously introduced [34]. Anchor points are calculated using a *hit-and-extend* approach where high-scoring spaced-word matches are used as *seeds*: for an underlying binary pattern of length ℓ representing *match* and *don't care* positions, we rapidly identify *spaced-word matches*, *i.e.* length- ℓ segment pairs from the input sequences with matching nucleotides at the *match* positions but with possible mismatches at the *don't care* positions. For each spaced-word match, we then calculate a similarity score considering *all* aligned positions – including the *don't-care positions* –, and we keep only those spaced-word matches that have a score above a certain threshold. These segment pairs are then extended to locally-maximal gap-free alignments, similar as in *BLAST* [1]. To evaluate our anchoring approach, we used the *Mugsy* pipeline using our software in the initial step, to find anchor points. For closely related input sequences, the quality of the resulting alignments is comparable to the original version of *Mugsy* where exact word matches are used for anchoring. Our approach is far superior, however, if distal sequences are to be aligned, where most other alignment approaches either fail to produce alignments or require an unacceptable amount of time.

Through our web site, we provide the adapted *Mugsy* pipeline with our anchoring approach as a pipeline for genome-sequence alignment that can be readily installed. A standalone version of our *spaced-words* software is provided as well, such that developers can integrate it into their own sequence-analysis pipelines.

Filtered Spaced Word Matches

For a sequence S of length L over an alphabet Σ and $0 < i \leq L$, $S[i]$ denotes the i -th symbol of S . For integers $w \leq \ell$, a binary *pattern* P of length ℓ and weight w is a word over $\{0, 1\}$ of length ℓ such that there are exactly w indices i with $P[i] = 1$. These positions are called *match positions*, while positions i with $P[i] = 0$ are called *don't-care positions*. A *spaced word* with respect to a pattern P is a word w over $\Sigma \cup \{*\}$ where $*$ is a wildcard character not contained in Σ , and $w[k] = *$ holds if and only if k is a *don't-care position*, i.e. if $P[k] = 0$, see also [33, 29]. A spaced word w with respect to a pattern P occurs in a sequence S at position i if $S[i + k - 1] = w[k]$ for all match positions k of the pattern P .

For sequences S_1 and S_2 with lengths L_1 and L_2 , respectively and a pattern P of length ℓ , and $1 \leq i \leq L_1 - \ell + 1, 1 \leq j \leq L_2 - \ell + 1$, we say that there is *spaced-word match* between S_1 and S_2 at (i, j) with respect to P if the spaced words at i in S_1 and at j in S_2 are identical - in other words, if for all match positions k in P , one has

$$S_1[i + k - 1] = S_2[j + k - 1].$$

Below is a spaced-word match between two *DNA* sequences S_1 and S_2 at $(5, 2)$ with respect to the pattern $P = 1100101$:

$$\begin{array}{rcccccccccccc} S_1 : & G & C & T & G & T & A & T & A & C & G & T & C \\ S_2 : & & & & S & T & A & C & A & C & T & T & A & T \\ P : & & & & & & 1 & 1 & 0 & 0 & 1 & 0 & 1 \end{array}$$

Indeed, the spaced word ‘ $TA**C*T$ ’ occurs at positions 5 in S_1 and at position 2 in S_2 .

Herein, we propose to use spaced-word matches as a first step to calculate *anchor points* for pairwise alignment. We therefore need some criterion to distinguish between spaced-word matches representing *true homologies* and random *background* matches. In a previous paper, we used spaced-word matches to estimate phylogenetic distances between genomic sequences [34]. To this end, we first identified all spaced-word matches with respect to a given pattern P . To remove spurious random spaced-word matches, we applied a simple *filtering procedure*: using a nucleotide substitution matrix [13], we calculated for each spaced-word match the sum scores of all aligned pairs of nucleotides (including match and don't-care positions), and we removed all spaced-word matches with a score below zero.

A graphical representation of the spaced-word matches between two sequences shows that this procedure can clearly separate random spaced-word




Figure 1: *Spaced-words histogram* for a comparison of two bacterial genomes, *Phaeobacter gallaeciensis* 2.10 and *Rhodobacteriales bacterium* Y4I. All possible spaced-word matches with respect to a given binary pattern P are identified, and their scores are calculated as explained in the main text. The number of spaced-word matches with a score s is plotted against s . Two peaks are visible, an approximately normally distributed peak for background spaced-word matches, and a more complex peak for spaced-word matches representing homologies. With a cut-off value of zero, background and homologous spaced-word matches can be reliably separated.

matches from true homologies. If we plot for each possible score value s the number of spaced-word matches with score s , we obtain a bimodal distribution with one peak for random matches and a second peak for homologies. We call such a plot a *spaced-words histogram*. For simulated sequence pairs under a simple model of evolution, both peaks are normally distributed. For real-world sequences, the random peak is still normally distributed, but the ‘homologous’ peak is more complex, see Figure 1. Even so, using a cut-off value of zero can clearly distinguish between random matches and true homologies. More examples for *spaced-words histograms* are given in [34].

Our approach to find anchor points for pairwise genomic alignment is as follows. For given parameters ℓ and w , we first calculate a binary pattern

with length ℓ and weight (number of *match* positions) w using our recently developed software *rasbhari* [25]. We then identify all spaced-word matches with respect to P . To find homologies even for distantly related sequences, we use patterns with a low weight; by default, we use a weight of $w = 10$. On the other hand, we use a large number of *don't-care* positions, since this makes it easier to distinguish true homologies from random spaced-word matches. By default, we use a pattern length of $\ell = 110$, so our patterns contain 10 match positions and 100 don't-care positions; we use the following nucleotide substitution matrix described in [13]:

	<i>A</i>	<i>C</i>	<i>G</i>	<i>T</i>
<i>A</i>	91	-114	-31	-123
<i>C</i>		100	-125	-31
<i>G</i>			100	-114
<i>T</i>				91

Based on this matrix, we calculate the *score* of each spaced-word match as the sum of the substitution scores of all aligned pairs of nucleotides. We then discard all spaced-word matches with a score below zero.

Next, we extend the identified spaced-word matches in both directions without gaps. As the starting point for this extension, we do not use the full spaced-word matches, but their mid points. The reason for this is that, with our long patterns, even a high-scoring spaced-word match may not represent sequence homologies over its entire length. It often occurs that some part of a spaced-word aligns homologous nucleotides, but another part extends into non-homologous regions of the sequences. There is a high probability, however, that the mid point of a long, high-scoring spaced-word match is located within a region of true homology. Finally, we use the produced ‘extended’ gap-free alignments as anchor points for alignment.

Evaluation

To evaluate *Filtered Spaced Word Matches (FSWM)* and to compare it to the state-of-the-art approach to alignment anchoring, we used the *Mugsy* software system. As mentioned above, the original *Mugsy* uses *MUMmer* to find pairwise anchor points. We replaced *MUMmer* in the *Mugsy* pipeline by our *FSWM*-based anchor points and evaluated the resulting multiple alignments. In addition, we compared these alignments to alignments produced by the multiple genome aligner *Cactus* [44]. *Cactus* is known to be one of the best existing tools for multiple genome alignment; it performed excellent

in the *Alignathon* study [23]. To measure the performance of the compared methods, we used simulated genomic sequences as well as three sets of real genomes. To make *MUMmer* directly comparable to *FSWM*, we used a minimum length of 10 *nt* for maximum unique matches, corresponding to the default *weight* (sum of *match positions*) used in *Spaced Words*. Note that, by default, *MUMmer* uses a minimum length of 15 *nt*. With this default value, however, we obtained alignments of much lower quality. The *Cactus* tool was run with default values.

Simulated genome sequences

To simulate genomic sequences, we used the *Artificial Life Framework (ALF)* developed by Dalquen *et al.* [14]. *ALF* evolves gene sequences based on a probabilistic model along a randomly generated tree, starting with an ancestral gene. During this process evolutionary events are logged such that the *true* MSA is known for each simulated gene family. This *true* MSA can then be used as reference to assess the quality of automatically generated alignments.

We generated a series of 14 data sets, each containing 30 simulated ‘genomes’, with increasing mutation rates for the different data sets. For all other parameters in *ALF*, we used the default settings. In each data set, there are 750 simulated gene families such that one gene from each gene family is present in each of the 30 simulated genomes. Thus, each of the ‘genomes’ contains the same set of 750 genes. We varied the mutation rates between an average of 0.1013 substitutions per position for the first data set to an average of 0.8349 substitutions per position for the 14th data set. The maximal pairwise distances between all pairs of sequences within one data set ranges from 0.1640 for the first to 1.0923 for the 14th data set. The simulated genes have an average length of about 1500bp, summing up to a total size of about 32 MB per data set.

To assess the quality of the produced alignments, we calculated *recall* and *precision* values in the usual way. If, for one given data set, S is the set of all positions in the 30 simulated genomes, we denote by $A \subset S^2$ the set of all pairs of positions aligned by the alignment that is to be evaluated while $R \subset S^2$ denotes the set of all pairs of positions aligned in the reference alignment. *recall* and *precision* are then defined as

$$recall = \frac{|A \cap R|}{|R|}, \quad precision = \frac{|A \cap R|}{|A|} \quad (1)$$

The harmonic mean of *recall* and *precision* is called the *balanced F-score* and

is often used as an overall measure of accuracy; it is thus defined as

$$F_{score} = 2 \times \frac{precision \times recall}{precision + recall}$$

To estimate these three values, we used the tool *mafComparator* which was also used in the *Alignathon* study [23]. Since it is impractical to consider the entire set S^2 of pairs of positions of the test sequences, we sampled 10 million pairs of positions for each data set. This corresponds to the evaluation procedure used in *Alignathon*.

For the simulated sequence sets, their *precision* and *recall* values are shown in Figure 2. For data sets with smaller mutation rates, alignments obtained with *FSWM* are only slightly better than those obtained with *MUMmer*. However, if the mutation rate increases, our spaced-words approach substantially outperforms the original version of *Mugsy* where exact word matches are used to find anchor points. Not only more homologies are detected but also the *precision* is slightly higher if *Filtered Spaced Word Matches* is used instead of *MUMmer*.

Real-world genome sequences

For real-world genome families, it is usually not possible to calculate the *precision* of MSA programs because it is, in general, not known which sequence positions exactly are homologous to each other and which ones are not. If there are *core blocks* of the sequences for which the biologically correct alignment is known, at least the *recall* can be calculated for these core blocks. For most genome sequences, however, no such core blocks are available. To evaluate *Mugsy*, the authors of the program used the number of *core columns* of the produced alignments as a criterion for alignment quality [2]. Here, a *core column* is defined as a column that does not contain gaps, *i.e.* a column that aligns nucleotides from all of the input sequences. In addition, the authors of *Mugsy* used the *number of pairs of aligned positions* of the aligned sequences as an indicator of alignment quality. In this paper, we are using the same criteria to evaluate multiple alignments of real-world genomes.

As a first real-world example, we used a set of 29 *E.coli/Shigella* genomes that has already been used in the original *Mugsy* paper, see *supplementary material* for details; these sequences have also been used to evaluate alignment-free methods [26, 49, 41]. The total size of this data set is about 141 MB. As a second test set, we used another prokaryotic data set which consists of 32 complete *Roseobacter* genomes (details in the *supplementary*




Figure 2: *Recall* and *Precision* of *Mugsy* with anchor points from *Filtered Spaced Word Matches (FSWM)* and *MUMmer*, respectively, and of *Cactus* on simulated genomic sequences generated with *ALF*, see main text for details. *FSWM* was used with the default *weight* $w = 10$, *i.e.* with 10 *match positions* in the underlying pattern. In addition, we ran *FSWM* with $w = 8$.

material). This data set was used to assess the performance on more distantly related organisms than the *E.coli/Shigella* strains. The total size of these data set is about 135 MB. To test our approach on eukaryotic genomes, we used as a third test case a set of nine fungal genomes, namely *Coprinopsis*




Figure 3: *F-Score* of *Mugsy* with anchor points from *Filtered Spaced Word Matches* and *MUMmer*, respectively, and of *Cactus* on simulated genomic sequences generated with *ALF*.

cinerea, *Neurospora crassa*, *Aspergillus terreus*, *Aspergillus nidulans*, *Histoplasma capsulatum*, *Paracoccidioides brasiliensis*, *Saccharomyces cerevisiae*, *Schizosaccharomyces pombe* and *Ustilago maydis* (genbank accession numbers are given in the supplementary material). The total size of this third data set is about 253 MB. The results of *Mugsy* with *MUMmer* and *FSWM* for the three real-world data sets are shown in Table 1, together with the results obtained with *Cactus*. In addition to the number of *core columns* and the number of aligned pairs of positions, the table contains the number of *core Locally Collinear Blocks*, *i.e.* the number of *Locally Collinear Blocks* involving all of the input sequences, and the total number of *Locally Collinear Blocks* returned by the alignment programs.

Program run time

Table 2 reports the program run times of *Mugsy* with *FSWM*, *Mugsy* with *MUMmer* and *Cactus* on the above three real-world sequence sets. In addition, the table contains the run times for *FSWM* and *MUMmer* alone.

	# core LCBs	# aligned pairs	# core col.	# LCBs
29 <i>E.coli/Shigella</i> genomes				
<i>Mugsy</i> + <i>MUMmer</i>	539	1,61E+09	2,827,115	4,138
<i>Mugsy</i> + <i>FSWM</i>	664	1,63E+09	2,867,432	5,906
<i>Cactus</i>	20,163	1,48E+09	2,663,750	56,592
32 <i>Roseobacter</i> genomes				
<i>Mugsy</i> + <i>MUMmer</i>	39	3,63E+08	13,654	13,501
<i>Mugsy</i> + <i>FSWM</i>	859	7,15E+08	824,054	30,836
<i>Cactus</i>	5,984	4,95E+08	280,085	337,320
9 fungal genomes				
<i>Mugsy</i> + <i>MUMmer</i>	9	5,88E+06	2,097	4,252
<i>Mugsy</i> + <i>FSWM</i>	2,590	1,18E+08	718,176	89,555
<i>Cactus</i>	31,589	1,33E+08	828,680	848,242

Table 1: Multiple alignments of 29 *E.coli/Shigella* genomes, 32 *Roseobacter* genomes and 9 fungal genomes, calculated with *Mugsy* using anchor points from our *spaced-words* approach and from *MUMs*, respectively, and with *Cactus*. The first column contains the number of *core columns*, *i.e.* the number of columns in the multiple alignment that do not contain gaps; the second column contains the total number of aligned pairs of positions in the alignment. The third column contains the number of *core Locally Collinear Blocks (LCBs)* *i.e.* the number of *LCBs* that involve *all* of the aligned genomes (‘core LCBs’), while the last column contains the total number of *LCBs*.

1 Discussion

In this paper, we proposed a novel approach to calculate anchor points for genome alignment. Finding suitable anchor points is a critical step in all methods for genome alignment, since the selected anchor points determine which regions of the sequences can be aligned to each other in the final alignment. A sufficient number of anchor points is necessary to keep the search space and run time of the main alignment procedure manageable, so *sensitive* methods are needed to find anchor points. Wrongly selected anchor points, on the other hand, can seriously deteriorate the quality of the final alignments, so anchoring procedures must also be highly *specific*.

Earlier approaches to genomic alignment used exact word matches as anchor points [18, 28], since such matches can be easily found using suffix trees and related indexing structures. These approaches are limited, however, to

	<i>E.coli/Shigella</i>	<i>Roseobacter</i>	<i>fungus genomes</i>
<i>FSWM</i>	59	83	110
<i>FSWM + Mugsy</i>	638	6428	1488
<i>MUMmer</i>	73	63	43
<i>MUMmer + Mugsy</i>	286	1099	63
<i>Cactus</i>	714	1775	775

Table 2: Run time in minutes for three different multiple genome-alignment methods applied to the three test data sets that we used in our program evaluation.

situations where closely related genomes are to be aligned, for example different strains of a bacterium. In modern approaches to database searching, *spaced seeds* are used to find potential sequence homologies [35, 27, 11]. Here, binary patterns of *match* and *don't care* positions are used, and two sequence segments of the corresponding length are considered to match if identical residues are aligned at the *match* positions, while mismatches are allowed at the *don't care* positions. Such pattern-based approaches are more *sensitive* than previous methods that relied on exact word matches.

We previously proposed to apply the ‘spaced-seeds’ idea to alignment-free sequence comparison, by replacing contiguous words by so-called *spaced words*, *i.e.* by words that contain wildcard characters at certain pre-defined positions [33]. More recently, we introduced *filtered spaced word matches* [34] to estimate phylogenetic distances between genome sequences. In the latter approach, we first identify spaced-word matches using relatively long patterns with only few *match* positions. For the identified matching segments, we then look at *all* aligned pairs of nucleotides, including the ones at the *don't-care* positions, and we discard spaced-word matches if the overall degree of similarity between the two segments is below a threshold. Phylogenetic distances can be estimated based on the aligned nucleotides at the *don't-care* positions of the remaining spaced-word matches. We showed that this procedure is fast and highly sensitive, and it can reliably distinguish between true homologies and spurious sequence similarities.

In the present study, we used filtered spaced word matches to calculate high-quality anchor points for genomic sequence alignment. Instead of using spaced-word matches directly as anchor points, we extend them into both directions, similar to the *hit-and-extend* approach to database searching. To evaluate these anchor points, we integrated them into the popular genome-alignment pipeline *Mugsy*. Test runs on simulated genome sequences show that, for closely related sequences, *Mugsy* produces alignments of high qual-

ity with both types of anchor points. For more distantly related sequences, however, the *recall* values of the program drop dramatically if anchor points are calculated with *MUMmer* while, with our spaced-word matches, one observes recall values close to 100% for distances up to around 0.7 substitutions per position.

For real-world genomes, it is more difficult to evaluate the performance of genome aligners since there is only limited information available on which positions are homologous to each other and which ones are not. Angiuoli and Salzberg [2] therefore used the number of aligned pairs of positions as an indicator of alignment quality, together with the size of the ‘core alignment’, *i.e.* the number of alignments columns that do not contain gaps. At first glance, these criteria might seem questionable; it would be trivial to maximize these values, simply by aligning sequences without internal gaps, by adding gaps only at the ends of the shorter sequences. However, as shown in Figure 2, all MSA programs in our study have high *precision* values, *i.e.* positions aligned by these programs are likely to be true homologs. In this situation, the number of aligned position pairs and size of the ‘core alignment’ can be considered as a proxy for the *recall* of the applied methods *i.e.* the proportion of homologies that are correctly aligned.

For distantly related sequence sets, the total run time of *Mugsy* is much higher with our *FSWM* anchoring approach than with *MUMmer*. One reason for the increased run time with *FSWM* is the fact that, with spaced-words, far more *Locally Collinear Blocks* are detected, than if exact word matches are used as anchor points, especially for distantly related sequences where exact word matching is not very sensitive. One possible solution for this issue would be to apply user-defined threshold values for the total number of returned *Locally Collinear Blocks* or for their similarity scores, to reduce the run time of the final alignment procedure for large genomic sequences.

References

- [1] S. F. Altschul, W. Gish, W. Miller, E. M. Myers, and D. J. Lipman. Basic local alignment search tool. *Journal of Molecular Biology*, 215:403–410, 1990.
- [2] S. V. Angiuoli and S. L. Salzberg. Mugsy: fast multiple alignment of closely related whole genomes. *Bioinformatics*, 27:334–342, 2011.

- [3] L. M. Barton, B. Göttgens, M. Gering, J. G. Gilbert, D. Grafham, J. Rogers, D. Bentley, R. Patient, and A. R. Green. Regulation of the stem cell leukemia (SCL) gene: a tale of two fishes. *Proc. Natl. Acad. Sci. USA*, 98:6747–6752, 2001.
- [4] S. Batzoglou. The many faces of sequence alignment. *Briefings in Bioinformatics*, 6:6–22, 2005.
- [5] M. Blanchette, W. J. Kent, C. Riemer, L. Elnitski, A. F. A. Smit, K. M. Roskin, R. Baertsch, K. Rosenbloom, H. Clawson, E. D. Green, D. Haussler, and W. Miller. Aligning multiple genomic sequences with the threaded blockset aligner. *Genome Research*, 14:708–715, 2004.
- [6] R. K. Bradley, A. Roberts, M. Smoot, S. Juvekar, J. Do, C. Dewey, I. Holmes, and L. Pachter. Fast statistical alignment. *PLOS Comput Biol*, 5:e1000392, 2009.
- [7] N. Bray, I. Dubchak, and L. Pachter. AVID: A global alignment program. *Genome Research*, 13:97–102, 2003.
- [8] N. Bray and L. Pachter. MAVID multiple alignment server. *Nucleic Acids Research*, 31:3525–3526, 2003.
- [9] M. Brudno, M. Chapman, B. Göttgens, S. Batzoglou, and B. Morgenstern. Fast and sensitive multiple alignment of large genomic sequences. *BMC Bioinformatics*, 4:66, 2003.
- [10] M. Brudno, C. Do, G. Cooper, M. Kim, E. Davydov, NISC Sequencing Consortium, E. Green, A. Sidow, and S. Batzoglou. LAGAN and multi-LAGAN: Efficient tools for large-scale multiple alignment of genomic DNA. *Genome Research*, 13:721–731, 2003.
- [11] B. Buchfink, C. Xie, and D. H. Huson. Fast and sensitive protein alignment using DIAMOND. *Nature Methods*, 12:59–60, 2015.
- [12] M. A. Chapman, F. J. Charchar, S. Kinston, C. P. Bird, D. Grafham, J. Rogers, F. Grützner, J. A. M. Graves, A. R. Green, and B. Göttgens. Comparative and functional analysis of LYL1 loci establish marsupial sequences as a model for phylogenetic footprinting. *Genomics*, 81:249–259, 2003.
- [13] F. Chiaromonte, V. B. Yap, and W. Miller. Scoring pairwise genomic sequence alignments. In R. B. Altman, A. K. Dunker, L. Hunter, and

- T. E. Klein, editors, *Pacific Symposium on Biocomputing*, pages 115–126, 2002.
- [14] D. A. Dalquen, M. Anisimova, G. H. Gonnet, and C. Dessimoz. ALF - a simulation framework for genome evolution. *Molecular Biology and Evolution*, 29:1115–1123, 2012.
 - [15] A. C. E. Darling, B. Mau, F. R. Blattner, and N. T. Perna. Mauve: multiple alignment of conserved genomic sequence with rearrangements. *Genome Research*, 14:1394–1403, 2004.
 - [16] A. E. Darling, B. Mau, and N. T. Perna. progressiveMauve: Multiple Genome Alignment with Gene Gain, Loss and Rearrangement. *PLOS ONE*, 5:e11147+, 2010.
 - [17] A. E. Darling, T. J. Treangen, L. Zhang, C. Kuiken, X. Messeguer, and N. T. Perna. *Algorithms in Bioinformatics: 6th International Workshop, WABI 2006, Zurich, Switzerland, September 11-13, 2006. Proceedings*, chapter Procrastination Leads to Efficient Filtration for Local Multiple Alignment, pages 126–137. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006.
 - [18] A. L. Delcher, S. Kasif, R. D. Fleischmann, J. Peterson, O. White, and S. L. Salzberg. Alignment of whole genomes. *Nucleic Acids Research*, 27:2369–2376, 1999.
 - [19] C. N. Dewey and L. Pachter. Evolution at the nucleotide level: the problem of multiple whole-genome alignment. *Human Molecular Genetics*, 15:R51–R56, 2006.
 - [20] A. Döring, D. Weese, T. Rausch, and K. Reinert. SeqAn – an efficient, generic C++ library for sequence analysis. *BMC Bioinformatics*, 9:11, 2008.
 - [21] I. Dubchak, A. Poliakov, A. Kislyuk, and M. Brudno. Multiple whole-genome alignments without a reference organism. *Genome Research*, 19:682–689, 2009.
 - [22] R. Durbin, S. R. Eddy, A. Krogh, and G. Mitchison. *Biological sequence analysis*. Cambridge University Press, Cambridge, UK, 1998.
 - [23] D. Earl, N. Nguyen, G. Hickey, R. S. Harris, S. Fitzgerald, K. Beal, I. Seledtsov, V. Molodtsov, B. J. Raney, H. Clawson, J. Kim, C. Kementa, J.-M. Chang, I. Erb, A. Poliakov, M. Hou, J. Herrero, W. J.

- Kent, V. Solovyev, A. E. Darling, J. Ma, C. Notredame, M. Brudno, I. Dubchak, D. Haussler, and B. Paten. Alignathon: a competitive assessment of whole-genome alignment methods. *Genome Research*, 24:2077–2089, 2014.
- [24] O. Gotoh. An improved algorithm for matching biological sequences. *J. Mol. Biol.*, 162:705–708, 1982.
- [25] L. Hahn, C.-A. Leimeister, R. Ounit, S. Lonardi, and B. Morgenstern. *rasbhari*: optimizing spaced seeds for database searching, read mapping and alignment-free sequence comparison. *PLOS Computational Biology*, 12(10):e1005107, 2016.
- [26] B. Haubold, F. Klötzl, and P. Pfaffelhuber. *andi*: Fast and accurate estimation of evolutionary distances between closely related genomes. *Bioinformatics*, 31:1169–1175, 2015.
- [27] H. Hauswedell, J. Singer, and K. Reinert. *Lambda*: the local aligner for massive biological data. *Bioinformatics*, 30:i349–i355, 2014.
- [28] M. Höhl, S. Kurtz, and E. Ohlebusch. Efficient multiple genome alignment. *Bioinformatics*, 18:312S–320S, 2002.
- [29] S. Horwege, S. Lindner, M. Boden, K. Hatje, M. Kollmar, C.-A. Leimeister, and B. Morgenstern. *Spaced words* and *kmacs*: fast alignment-free sequence comparison based on inexact word matches. *Nucleic Acids Research*, 42:W7–W11, 2014.
- [30] W. Huang, D. M. Umbach, and L. Li. Accurate anchoring alignment of divergent sequences. *Bioinformatics*, 22:29–34, 2006.
- [31] S. Kurtz. Reducing the space requirement of suffix trees. *Software – Practice and Experience*, 29:1149–1171, 1999.
- [32] S. Kurtz, A. Phillippy, A. L. Delcher, M. Smoot, M. Shumway, C. Antonescu, and S. L. Salzberg. Versatile and open software for comparing large genomes. *Genome Biology*, 5:R12+, 2004.
- [33] C.-A. Leimeister, M. Boden, S. Horwege, S. Lindner, and B. Morgenstern. Fast alignment-free sequence comparison using spaced-word frequencies. *Bioinformatics*, 30:1991–1999, 2014.
- [34] C.-A. Leimeister, S. Sohrabi-Jahromi, and B. Morgenstern. Fast and accurate phylogeny reconstruction using filtered spaced-word matches. *Bioinformatics*, 10.1093/bioinformatics/btw776 (in press).

- [35] M. Li, B. Ma, D. Kisman, and J. Tromp. PatternHunter II: Highly sensitive and fast homology search. *Genome Informatics*, 14:164–175, 2003.
- [36] B. Ma, J. Tromp, and M. Li. PatternHunter: faster and more sensitive homology search. *Bioinformatics*, 18:440–445, 2002.
- [37] B. Morgenstern. A simple and space-efficient fragment-chaining algorithm for alignment of DNA and protein sequences. *Applied Mathematics Letters*, 15:11–16, 2002.
- [38] B. Morgenstern, A. Dress, and T. Werner. Multiple DNA and protein sequence alignment based on segment-to-segment comparison. *Proceedings of the National Academy of Sciences*, 93:12098–12103, 1996.
- [39] B. Morgenstern, S. J. Prohaska, D. Pöhler, and P. F. Stadler. Multiple sequence alignment with user-defined anchor points. *Algorithms for Molecular Biology*, 1:6, 2006.
- [40] B. Morgenstern, O. Rinner, S. Abdeddaïm, D. Haase, K. Mayer, A. Dress, and H.-W. Mewes. Exon discovery by genomic sequence alignment. *Bioinformatics*, 18:777–787, 2002.
- [41] B. Morgenstern, B. Zhu, S. Horwege, and C.-A. Leimeister. Estimating evolutionary distances between genomic sequences from spaced-word matches. *Algorithms for Molecular Biology*, 10:5, 2015.
- [42] S. B. Needleman and C. D. Wunsch. A general method applicable to the search for similarities in the amino acid sequence of two proteins. *J. Mol. Biol.*, 48:443–453, 1970.
- [43] A. Y. Ogurtsov, M. A. Roytberg, S. A. Shabalina, and A. S. Kondrashov. OWEN: aligning long collinear regions of genomes. *Bioinformatics*, 18:1703–1704, 2002.
- [44] B. Paten, D. Earl, N. Nguyen, M. Diekhans, D. Zerbino, and D. Hausler. Cactus: Algorithms for genome multiple sequence alignment. *Genome Research*, 21:1512–1528, 2011.
- [45] B. Raphael, D. Zhi, H. Tang, and P. Pevzner. A novel method for multiple alignment of sequences with repeated and shuffled elements. *Genome Research*, 14:2336 – 2346, 2004.

- [46] T. Rausch, A.-K. Emde, D. Weese, A. Döring, C. Notredame, and K. Reinert. Segment-based multiple sequence alignment. *Bioinformatics*, 24:i187–i192, 2008.
- [47] T. F. Smith and M. S. Waterman. Identification of common molecular subsequences. *Journal of Molecular Biology*, 147:195–197, 1981.
- [48] The Arabidopsis Genome Initiative. Analysis of the genome sequence of the flowering plant *Arabidopsis thaliana*. *Nature*, 408:796–815, 2000.
- [49] H. Yi and L. Jin. Co-phylog: an assembly-free phylogenomic approach for closely related organisms. *Nucleic Acids Research*, 41:e75, 2013.